arXiv:1602.08297v1 [q-fin.PM] 26 Feb 2016

Fluctuation-bias trade-off in portfolio optimization under Expected Shortfall with ℓ_2 regularization

Gábor Papp[1], Fabio Caccioli[2,3] and Imre Kondor[4,5]
1- Eötvös Loránd University, Institute for Physics, Budapest, Hungary
2- University College London, Department of Computer Science, London, WC1E 6BT, UK
3- Systemic Risk Centre, London School of Economics and Political Sciences, London, UK
4-Parmenides Foundation, Pullach, Germany
5- Department of Investment and Corporate Finance, Corvinus University of Budapest, Budapest, Hungary

Abstract

The optimization of a large random portfolio under the Expected Shortfall risk measure with an ℓ_2 regularizer is carried out by analytical calculation. The regularizer reins in the large sample fluctuations and the concomitant divergent estimation error, and eliminates the phase transition where this error would otherwise blow up. In the data-dominated region, where the number of different assets in the portfolio is much less than the length of the available time series, the regularizer plays a negligible role, while in the (much more frequently occurring in practice) opposite limit, where the samples are comparable or even small compared to the number of different assets, the optimum is almost entirely determined by the regularizer. Our results clearly show that the transition region between these two extremes is relatively narrow, and it is only here that one can meaningfully speak of a trade-off between fluctuations and bias.

1 Introduction

The optimization of large institutional portfolios belongs to the realm of high-dimensional statistics [1]: the dimension N, the number of different assets (or risk factors), is typically large compared to the sample size T (the length of the available time series). This poses a serious problem to optimization under any risk measure, which becomes particularly acute in the case of Value at Risk (VaR) and Expected Shortfall (ES), which discard a large part of the data except those at or above a high quantile. Without some kind of regularization this leads to a phase transition: above a critical value of the ratio $r = N/T$ (which depends on the confidence level α that determines the VaR threshold) it is impossible to carry out

the optimization, and upon approaching this critical value from below the estimation error increases without bound.

The estimation error problem of portfolio selection has been the subject of a large number of works [2–12]. The critical behavior and the locus of the phase boundary separating the region where the optimization is feasible from the one where it is not has also been studied in a series of papers [13–17].

In the present note we wish to discuss the effect of adding an ℓ_2 regularizer to the ES risk measure. As noted in [18] and [19], the optimization problem so obtained corresponds to support vector regression [20], therefore its study is of interest also for machine learning, independently of the portfolio optimization context. Concerning its specific financial application, ℓ_2 regularization may have two different interpretations. First, ℓ_2 has the tendency of pulling the solution towards the diagonal where all the weights are the same. As such, ℓ_2 represents a diversification pressure [21,22] that may be useful in a situation where e.g. the market is dominated by a small number of blue chips. Alternatively, the portfolio manager may wish to take into account the market impact of the future liquidation of the portfolio already at its construction. As shown in [17], market impact considerations lead one to regularized portfolio optimization, with ℓ_2 corresponding to linear impact.

In this paper we carry out the optimization of ℓ_2-regularized ES analytically in the special case of i.i.d. Gaussian distributed returns, in the limit where both the dimension and the sample size are large, but their ratio $r = N/T$ is fixed. The calculation will be performed by the method of replicas borrowed from the statistical physics of disordered systems [23]. The present work extends a previous paper [24] by incorporating the regularizer. As expected, by preventing the phase transition from taking place, the regularizer fundamentally alters the overall picture (in this respect, the role of the regularizer is analogous to that of an external field coupled to the order parameter in a phase transition). As the technical details of the replica calculation have been laid out in [13] and, in a somewhat different form, in [17], we omit them in here. Instead we just recall the setup of the problem and display the solutions for the most important quantities: the relative estimation error of ES, and the susceptibility, i.e. the sensitivity of the optimal portfolio weights to small shifts in the returns. Our results exhibit a clear distinction between the region in the space of parameters where data dominate and the regularizer plays a minor role, from the one where the data are insufficient and the regularizer stabilizes the estimate at the price of essentially suppressing the data. Thereby, our results provide a clean and explicit example of the estimation error-bias trade-off that is at the heart of the regularization procedure. We find that the transition between the data-dominated regime and the bias-dominated one is rather sharp, so the region where an actual trade-off can take place is narrow. We also note that the limits of such a narrow transition range might be hard to identify on the basis of empirical data, but show up very clearly in our analytical approach.

2 The optimization of regularized ES

The optimization of ℓ_2-regularized ES can be formulated as [17]:

$$\min_{\vec{w},\vec{u},\epsilon} \left[(1-\alpha)T\epsilon + \sum_{t=1}^{T} u_t + \eta \sum_i w_i^2 \right], \tag{1}$$

$$\text{s.t.} \quad \vec{w}\cdot\vec{x}_t + \epsilon + u_t \geq 0; \quad u_t \geq 0; \quad \forall t, \tag{2}$$

$$\sum_i w_i = N, \tag{3}$$

where the variables w's are the portfolio weights, ϵ and $\vec{u}$ are auxiliary variables, x_{it} are i.i.d. normally distributed returns, and η is the amplitude of the regularizer. The mapping of the portfolio optimization problem under unregularized ES onto a linear program was established by [25]. The difference with respect to the original formulation is the term proportional to η that appears in the cost function, which is due to the ℓ_2 regularized.

Note that the sum of the weights w_i is normalized to N instead of the usual 1, in order to make the optimal weights of order unity. Also, the constraint on the expected return has been omitted for simplicity, so we are seeking the global optimum of the portfolio here. If the returns x_{it} are i.i.d. Gaussian variables and $N, T \to \infty$ with $r = N/T$ fixed, the method of replicas allows one to reduce the above optimization task in $N+T+1$ variables to the optimization of a cost function depending on only six variables, the so-called order parameters [13, 17]:

$$\begin{aligned} F(\lambda,\epsilon,q_0,\Delta,\hat{q}_0,\hat{\Delta}) &= \lambda + \frac{1}{r}(1-\alpha)\epsilon - \Delta\hat{q}_0 - \hat{\Delta}q_0 \\ &+ \langle \min_w [V(w,z)] \rangle_z + \frac{\Delta}{2r\sqrt{\pi}} \int_{-\infty}^{\infty} ds e^{-s^2} g\left(\frac{\epsilon}{\Delta} + s\sqrt{\frac{2q_0}{\Delta^2}} \right), \end{aligned} \tag{4}$$

where

$$V(w,z) = \hat{\Delta}w^2 - \lambda w - zw\sqrt{-2\hat{q}_0} + \eta w^2, \tag{5}$$

$\langle\cdot\rangle$ is an average over the standard normal variable z, and

$$g(x) = \begin{cases} 0, & x \geq 0 \\ x^2, & -1 \leq x \leq 0 \\ -2x-1, & x < -1 \end{cases} . \tag{6}$$

The stationarity conditions are readily derived to be:

$$1 = \langle w^* \rangle_z \tag{7}$$

$$(1-\alpha)+\frac{1}{2\sqrt{\pi}}\int_{-\infty}^{\infty}dse^{-s^2}g'\left(\frac{\epsilon}{\Delta}+s\sqrt{\frac{2q_0}{\Delta^2}}\right)=0 \tag{8}$$

$$\hat{\Delta}-\frac{1}{2r\sqrt{2\pi q_0}}\int_{-\infty}^{\infty}dse^{-s^2}sg'\left(\frac{\epsilon}{\Delta}+s\sqrt{\frac{2q_0}{\Delta^2}}\right)=0 \tag{9}$$

$$-\hat{q}_0-2\frac{\hat{\Delta}q_0}{\Delta}+\frac{1}{2r\sqrt{\pi}}\int_{-\infty}^{\infty}dse^{-s^2}g\left(\frac{\epsilon}{\Delta}+s\sqrt{\frac{2q_0}{\Delta^2}}\right)+\frac{(1-\alpha)}{r}\frac{\epsilon}{\Delta}=0 \tag{10}$$

$$\Delta=\frac{1}{\sqrt{-2\hat{q}_0}}\left\langle w^*z\right\rangle_z \tag{11}$$

$$q_0=\left\langle w^{*2}\right\rangle_z, \tag{12}$$

where the variable w^* is that value of the weight that minimizes the "potential" V in (5).

Three of the order parameters are easily eliminated and the integrals can be reduced to the error function and its integrals by repeated integration by parts, as in [24]. Finally, one ends up with three equations to be solved:

$$r\left(1-2\eta\Delta\right)=\Phi\left(\frac{\Delta+\epsilon}{\sqrt{q_0}}\right)-\Phi\left(\frac{\Delta+\epsilon}{\sqrt{q_0}}\right) \tag{13}$$

$$\alpha=\frac{\sqrt{q_0}}{\Delta}\left\{\Psi\left(\frac{\Delta+\epsilon}{\sqrt{q_0}}\right)-\Psi\left(\frac{\Delta+\epsilon}{\sqrt{q_0}}\right)\right\} \tag{14}$$

$$\frac{1}{2\Delta^2}+\frac{\alpha}{r}\frac{\epsilon}{\Delta}+\frac{q_0}{2\Delta^2}+\frac{1}{2r}-\frac{2\eta q_0}{\Delta^2}=\frac{q_0}{r\Delta^2}\left\{W\left(\frac{\Delta+\epsilon}{\sqrt{q_0}}\right)-W\left(\frac{\Delta+\epsilon}{\sqrt{q_0}}\right)\right\}, \tag{15}$$

where

$$\Phi(x)=\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x}dte^{-t^2/2}, \tag{16}$$

$$\Psi(x)=x\Phi(x)+\frac{1}{\sqrt{2\pi}}e^{-x^2/2}, \tag{17}$$

$$W(x)=\frac{x^2+1}{2}\Phi(x)+\frac{x}{2}\frac{1}{\sqrt{2\pi}}e^{-x^2/2}. \tag{18}$$

These functions are closely related to each other:

$\Psi'(x) = \Phi(x)$, and $W'(x) = \Psi(x)$.

As explained in [24], each of the three remaining order parameters in the above set of equations, q_0, Δ, and ϵ has a direct financial meaning: If $\mathrm{ES}_{\mathrm{out}}$ is the out-of-sample estimate of ES based on samples of size T, and $\mathrm{ES}^{(0)}$ is its true value (that would obtain for N finite and $T \to \infty$), then

$$\frac{\mathrm{ES}_{out}}{\mathrm{ES}^{(0)}} = \sqrt{\langle w^2 \rangle} = \sqrt{1 + \sigma_w^2} = \sqrt{q_0}, \tag{19}$$

that is $\sqrt{q_0} - 1$ is the relative estimation error of the out-of-sample estimate. Furthermore, the sensitivity of the sample average of the optimal portfolio weights to a small shift ξ in the returns is given by the variable Δ:

$$\left. \frac{\partial \langle w^* \rangle_z}{\partial \xi} \right|_{\xi=0} = \Delta. \tag{20}$$

We will call the variable Δ susceptibility in the following.

Finally, the variable ϵ is the in-sample estimate of Value at Risk.

The task is now to solve the stationary conditions and get the cost function by substituting the solutions back into eq. (4). The in-sample value of Expected Shortfall is simply related to the cost function as:

$$ES = \frac{rF}{1 - \alpha}. \tag{21}$$

The fundamental cause of the divergence of estimation error in the original, non-regularized problem was that ES as a risk measure is not bounded from below. In finite samples it can happen that one of the assets, or a combination of assets, dominates the others (i.e. produces a larger return than the others in the given sample) thereby leading to an apparent arbitrage: one can achieve an arbitrarily large gain (an arbitrarily large negative ES) by going very long in the dominant asset and correspondingly short in the others [26]. It is evident that the apparent arbitrage is a mere statistical fluctuation, but along a special curve in the $r - \alpha$ plane this divergence occurs with probability one [13]. As a result, the estimation error and susceptibility will diverge along this phase boundary. The phase boundary is shown in Fig. 1.

The role of regularization is to penalize the large fluctuations of the weight vector, thereby eliminating this phase transition. Since ES is a linear function of the weights, the quadratic regularizer will overcome the excessive fluctuations however small the coefficient η. Deep

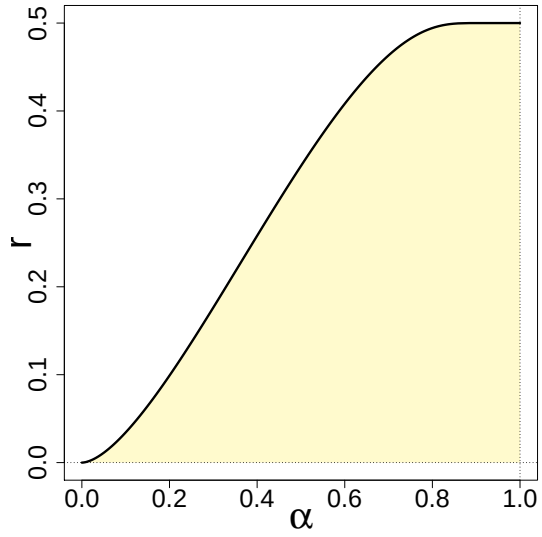

Figure 1: The phase boundary of unregularized ES for i.i.d. normal underlying returns. In the region below the phase boundary the optimization of ES is feasible and the estimation error is finite. Approaching the phase boundary from below, the estimation error diverges, and above the line optimization is no longer feasible.

inside the region of stability (shown by pale yellow in Fig. 1), a weak regularizer (small η) will modify the behavior of various quantities very little. In contrast, close to the phase boundary, and especially in the vicinity of the point $\alpha = 1$, $r = 0.5$ where the solution has an essential singularity, the effect of even a small η is very strong, and beyond the yellow region where originally there was no solution, the regularizer will dominate the scene. In the region where the solutions are stable even without the regularizer, $r = N/T$ is small, so we have an abundance of data. We call this region the data-dominated region. In the presence of the regularizer we will find finite solutions also far beyond the phase boundary, but here the role of the regularizer is decisive, so we can call this domain the bias-dominated region.

3 Results

The solution of the stationarity conditions can be obtained with the help of a computer. In the following, we present the numerical solutions for the relative estimation error and the susceptibility. The results will be displayed by constructing the contour maps of these quantities, which will allow us to make a direct comparison between our present results and those in [24].

In Fig. 2 we recall the contour map of the relative estimation error of ES without regular-

ization.

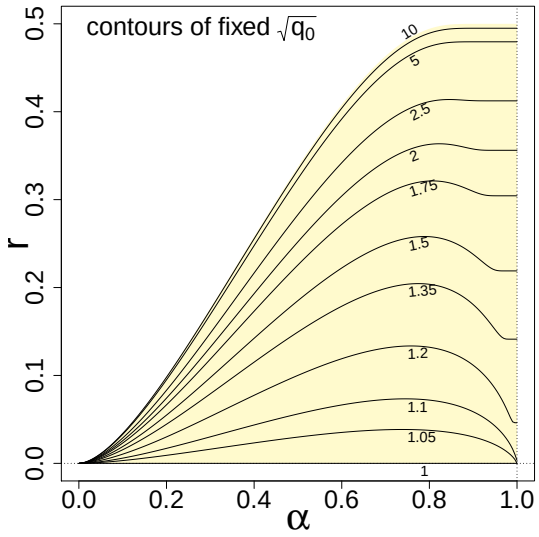

Figure 2: Contour lines of fixed $\sqrt{q_0}$ in the absence of regularization. These curves are also the contour lines for the relative error for the out-of-sample estimate of ES.

As can be seen, without regularization the constant q_0 curves are all inside the feasible region. For larger and larger values of q_0 the corresponding curves run closer and closer to the phase boundary, along which the estimation error diverges. Note that the phase boundary becomes flat, with all its derivatives vanishing, at the upper right corner of the feasible region: the point $\alpha = 1, r = 0.5$ is the locus of an essential singularity.

The estimation error problem is very clearly illustrated in this figure: the curves corresponding to an acceptably small relative error are the lowest ones among the q_0 contour lines, and the value of $r = N/T$ corresponding to a confidence limit α in the vicinity of 1 (such as the regulatory value $\alpha = 0.975$) are extremely small on these low lying curves. These small values of r require an unrealistically large sample T if N is not small. For example, at the regulatory value $\alpha = 0.975$, to achieve an estimation error smaller than 5%, for a portfolio with $N = 100$ assets one would need a time series of more than 7200 data points [24].

Let us see how regularization reorganizes the set of constant q_0 curves. Figs. 3 and 4. display these curves for two different values of the amplitude η of the regularizer. (Notice the logarithmic scale on the vertical axes in these figures.)

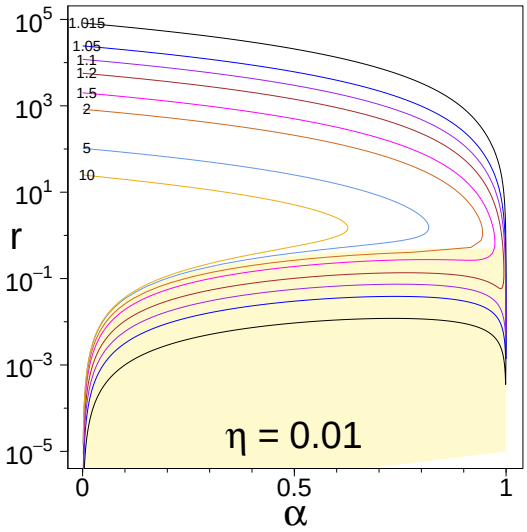

Figure 3: Contour plot for fixed values of $\sqrt{q_0}$ on the $\alpha - r$ plane at $\eta = 0.01$.

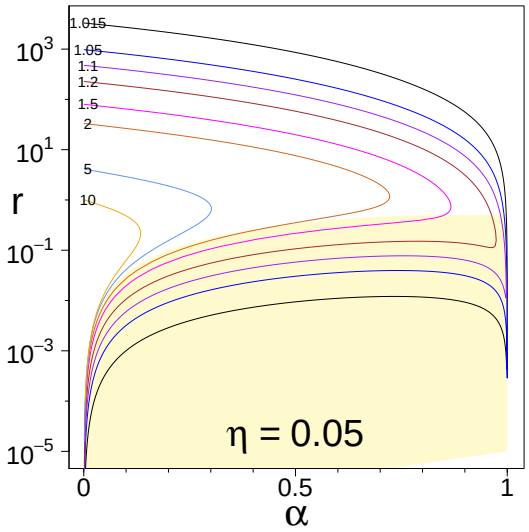

Figure 4: Contour plot for fixed values of $\sqrt{q_0}$ on the $\alpha - r$ plane at $\eta = 0.05$.

The curves of constant q_0 have two branches now. For a given q_0 the lower branch lies mostly or partly in the previously feasible region, the upper branch lies outside, above it. Along the lower branch the value of the ratio r is small, that means we have very large samples with respect to the dimension: this is the data-dominated regime. We can also see that when the data dominate, the dependence on the regularizer is weak: the set of curves inside the yellow region is quite similar in the two figures, even though the amplitude of the regularizer has been increased 5-fold from Fig. 3 to Fig. 4. Following the curve corresponding to a given value of q_0, say the black one, we see that at the beginning it is increasing with α, then in the vicinity of $\alpha = 1$ it starts to decline, then it sharply turns around and shoots up steeply, leaving the feasible region and increasing with decreasing α. Along this upper branch the ratio r is not small any more. We do not have large samples here, in fact, the situation is just the opposite: the dimension N becomes larger than the size T of the samples. Clearly, in this regime the regularizer dominates, this is the bias-dominated regime. It is important to note the sudden turn over of the curve in the vicinity of the original phase boundary. Such a sharp feature would be extremely hard to discover if we wanted to solve the original optimization problem by numerical simulations: the simulation would jump over to the upper branch before we could observe the sharp dip. This is even more so for real life data which are inevitably noisy.

Fig.5. shows the contour map of the susceptibility Δ without regularization. The curves of reasonably small values of the susceptibility lie low again, corresponding to small r's, i.e. very large samples. In Fig. 6 we exhibit one of these contour lines (the one corresponding

to $\Delta = 1$) as it is shifted upwards with increasing η. As can be seen from the figure, even a relatively small value of the regularizer can shift these lines significantly upwards.

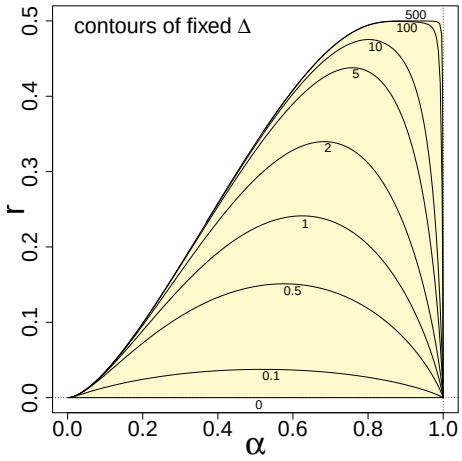

Figure 5: Contour lines of fixed Δ, the susceptibility of the portfolio weights to small shifts in the returns, with no regularization.

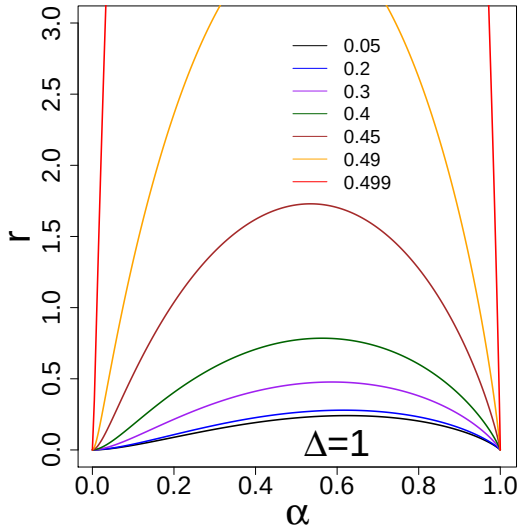

Figure 6: The lines of constant Δ are shifted upwards by regularization. The figure displays this on the example of the $\Delta = 1$ line for various values of the amplitude of the regularizer η. These values correspond to the different colors as shown in the figure.

Fig. 7 is the contour map of Δ with regularization. As can be seen, the iso-Δ lines are pushed upwards very strongly by regularization (the value of the amplitude of the

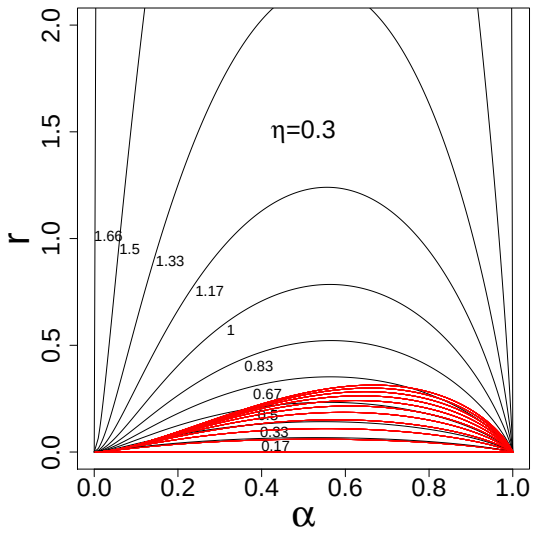

Figure 7: Contour map of Δ. Note how the lines of constant Δ are blown up by regularization. The amplitude of the regularizer is $\eta = 0.3$ here. The red curves lying within the originally allowed region are the unregularized Δ contour lines, corresponding to the same Δ values, as the black, regularized Δ lines.

regularizer is fixed at $\eta = 0.3$ in this figure). They grow far beyond the originally allowed region, which means the sensitivity of the optimal portfolio to changes in the returns gets diminished. This is explained by the fact that in this region it is the regularizer that determines the solution; the data and their small changes do not matter much here.

An important point in regularization is the correct choice of the parameter η. When data come from real observations, and the size of the sample (or the number of samples) is limited by time and/or cost considerations, the standard procedure is cross validation [27], i.e. using a part of the data to infer the value of η and checking the correctness of this choice on the other part. In the present analytical approach we have the luxury of infinitely large samples, so we can obtain the value of the amplitude of regularization by demanding a given relative error (that is a given q_0) for a given confidence limit α and given aspect ratio $r = N/T$. Fig. 8 displays the plot of r vs. η for $\alpha = 0.975$ as demanded by the new market risk regulation [28], and relative errors of 1%, 5% and 10%, respectively. It shows the change-over between the data-dominated resp. bias-dominated regimes very clearly. For a given value of r the corresponding value of η can be read off from the curves. If r is small (i.e. the sample is large with respect to the dimension) the curves with the prescribed values of relative error are almost horizontal. This means that when we have sufficient data the value of the regularizer is more or less immaterial: within reasonable limits we can choose any amplitude for the regularizer, it will not change the precision of the estimate, because in this situation the data will determine the optimum. Conversely, when the data are insufficient (r is not small, or it is even beyond the feasible region),

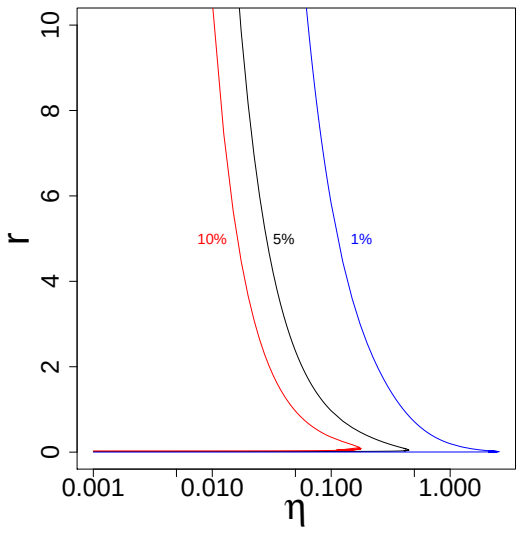

Figure 8: The value of $r = N/T$ vs. the coefficient of the regularizer for a fixed value of the confidence limit $\alpha = 0.975$ and three different values (1%, 5% and 10%) of the relative estimation error. The data-dominated and bias-dominated regions can be clearly distinguished in these curves: in the range of small r's the curves depend on the strength of the regularizer very weakly, while for r's in the vicinity of the phase boundary, and even more for large r's high in the originally unfeasible region, the fixed estimation error curves display a strong dependence on η. Conversely, for a given η we find two values of r, one small, the other large, corresponding to the same estimation error.

the value of η necessary to enforce a given relative error depends very much on r. In this region, however, we need a smaller and smaller η to find the same relative error, because here the data almost do not count and even a small regularizer will establish the optimum. The transition between these two regimes takes place around the points where the curves turn back. This happens still inside the feasible region, and the width of this range in r is rather small, always within a factor of 10. This is tantamount to saying that the range of sample sizes within which there is a real trade-off between fluctuation and bias is fairly narrow.

4 Conclusion

We have considered the problem of optimizing Expected Shortfall in the presence of an ℓ_2 regularizer. The regularizer takes care of the large sample fluctuations and eliminates the phase transition that would be present in the problem without regularization. Deep inside the feasible region where we have a large amount of data relative to the dimension, the size of the sample needed for a given level of relative estimation error is basically constant, largely independent of the regularizer. In the opposite case, for sample sizes compara-

ble to or small relative to the dimension, the regularizer is dominating the optimization and suppresses the data. The transition between the the data-dominated regime and the regularizer-dominated one is rather narrow, taking place roughly within a decade of sample sizes. It is in this transition region where we can meaningfully speak about a trade-off between fluctuation and bias, otherwise one or the other dominates the estimation. The identification of this transitional zone is very easy within the present scheme, where we could carry out the optimization analytically: the transitional zone is the small region where the curves in Fig. 8 sharply turn back. In real life, where the size of the samples can rarely be changed at will and where all kinds of external noises (other than those coming from the sample fluctuations) may be present, the distinction between the region where the data dominate and that where the bias does may be much less clear, and one may not be sure where the transition takes place between them. Below this transition there is not much point about using regularization, because the data themselves are sufficient to provide a stable and reliable estimate. Above the transition zone it is almost pointless to talk about the observed data, because they are crowded out by the bias. The result that the transition zone between these two extremes is narrow is the main message of this paper.

5 Acknowledgements

We are obliged to R. Kondor, M. Marsili and S. Still for valuable discussions. FC acknowledges support of the Economic and Social Research Council (ESRC) in funding the Systemic Risk Centre (ES/K002309/1). IK is grateful for the hospitality extended to him at the Computer Science Department of the University College of London where this paper has been written up.

References

[1] P. Bühlmann and S. Van De Geer. Statistics for high-dimensional data: methods, theory and applications. Springer Science & Business Media, 2011.

[2] P. Jorion. Bayes-Stein estimation for portfolio analysis. Journal of Financial and Quantitative Analysis, 21:279–292, 1986.

[3] P.A. Frost and J.E. Savarino. An empirical Bayes approach to efficient portfolio selection. Journal of Financial and Quantitative Analysis, 21:293–305, 1986.

[4] O. Ledoit and M. Wolf. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. Journal of Empirical Finance, 10(5):603–621, 2003.

[5] O. Ledoit and M. Wolf. A well-conditioned estimator for large-dimensional covariance matrices. J. Multivar. Anal., 88:365–411, 2004.

[6] O. Ledoit and M. Wolf. Honey, I shrunk the sample covariance matrix. J. Portfolio Management, 31:110, 2004.

[7] V. Golosnoy and Y. Okhrin. Multivariate shrinkage for optimal portfolio weights. The European Journal of Finance, 13:441–458, 2007.

[8] J. Brodie, I. Daubechies, C. De Mol, D. Giannone, and I. Loris. Sparse and stable Markowitz portfolios. Proceedings of the National Academy of Science, 106(30):12267–12272, 2009.

[9] L. Laloux, P. Cizeau, J.-P. Bouchaud, and M. Potters. Noise dressing of financial correlation matrices. Phys. Rev. Lett., 83:1467–1470, 1999.

[10] L. Laloux, P. Cizeau, J.-P. Bouchaud, and M. Potters. Random matrix theory and financial correlations. International Journal of Theoretical and Applied Finance, 3:391, 2000.

[11] V. Plerou, P. Gopikrishnanm, B. Rosenow, L.A.N. Amaral, and H.E. Stanley. Universal and non-universal properties of cross-correlations in financial time series. Phys. Rev. Lett., 83:1471, 1999.

[12] V. Plerou, P. Gopikrishnan, B. Rosenow, L. A. N. Amaral, T. Guhr, and H. E. Stanley. A random matrix approach to cross-correlations in financial time-series. Phys. Rev. E, 65:066136, 2000.

[13] S. Ciliberti, I. Kondor, and M. Mézard. On the feasibility of portfolio optimization under expected shortfall. Quantitative Finance, 7:389–396, 2007.

[14] I. Varga-Haszonits and I. Kondor. The instability of downside risk measures. J. Stat. Mech., P12007, 2008.

[15] S. Ciliberti and M. Mézard. Risk minimization through portfolio replication. Eur. Phys. J., B 57:175–180, 2007.

[16] F. Caccioli, S. Still, M. Marsili, and I. Kondor. Optimal liquidation strategies regularize portfolio selection. The European Journal of Finance, 19(6):554–571, 2013.

[17] F. Caccioli, I. Kondor, M. Marsili, and S. Still. Lp regularized portfolio optimization. Available at SSRN 2425326 and arXiv:1404.4040, 2014.

[18] A. Takeda and M. Sugiyama. ν-support vector machine as conditional value-at-risk minimization. In Proceedings of the 25th international conference on Machine learning, pages 1056–1063. ACM, 2008.

[19] S. Still and I. Kondor. Regularizing portfolio optimization. New Journal of Physics, 12(7):075034, 2010.

[20] F. Pérez-Cruz, J. Weston, D.J.L. Herrmann, and B. Scholkopf. Extension of the nu-svm range for classification. NATO SCIENCE SERIES SUB SERIES III COMPUTER AND SYSTEMS SCIENCES, 190:179–196, 2003.

[21] J.-P. Bouchaud and M. Potters. Theory of financial risk and derivative pricing. Cambridge Univ. Press, 2003.

[22] A. Gábor and I. Kondor. Portfolios with nonlinear constraints and spin glasses. Physica A: Statistical Mechanics and its Applications, 274(1):222–228, 1999.

[23] M. Mézard, G. Parisi, and M. A. Virasoro. Spin glass theory and beyond. World Scientific Lecture Notes in Physics Vol. 9, World Scientific, Singapore, 1987.

[24] F. Caccioli, I. Kondor, and G. Papp. Portfolio optimization under expected shortfall: contour maps of estimation error. arXiv preprint arXiv:1510.04943, 2015.

[25] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. Journal of Risk, 2(3):21–41, 2000.

[26] I. Kondor and I. Varga-Haszonits. Instability of portfolio optimization under coherent risk measures. Advances in Complex Systems, 13(03):425–437, 2010.

[27] T. Hastie, R. Tibshirani, and J. Friedman. The elements of statistical learning, data mining, inference, and prediction. Second edition. Springer series in statistics Springer, Berlin, 2008.

[28] Basel Committee on Banking Supervision. Minimum capital requirements for market risk. Bank for International Settlements, 2016.